arXiv:1405.6500v2 [cs.DB] 10 Jun 2014

Towards Efficient Path Query on Social Network with Hybrid RDF Management *

Lei Gai, Wei Chen **, Zhichao Xu, Changhe Qiu, and Tengjiao Wang

School of Electronic Engineering and Computer Science,
Peking University, Beijing, China
{lei.gai,pekingchenwei,zhchxu,chhqiu,tjwang}@pku.edu.cn

Abstract. The scalability and flexibility of Resource Description Framework(RDF) model make it ideally suited for representing Online Social Networks(OSN). One basic operation in OSN is to find chains of relations, such as k-Hop friends. Property path query in SPARQL can express this type of operation, but its implementation suffers from performance problem considering the ever growing data size and complexity of OSN. In this paper, we present a main memory/disk based hybrid RDF data management framework for efficient property path query. This hybrid framework realizes an efficient in-memory algebra operator for property path query using graph traversal, and estimates the cost of this operator to cooperate with existing cost-based optimization. Experiments on benchmark and real dataset demonstrated that our approach achieves a good tradeoff between data load expense and online query performance.

1 Introduction

In the age of Web 2.0, OSN have gained pervasive interests in both research communities and industries. There is a trend to model OSN using Semantic Web technologies, especially the vocabularies from FOAF[1] and SIOC[2] project. RDF, originally designed for the Semantic Web, have been wildly adopted for representing such kind of linked data. SPARQL[3] as the de-facto RDF query language is used for Basic Graph Pattern(BGP) query with bounded or unbounded variables. The scalable graph representation and flexible query capability make RDF model suited for large-scale complex OSN management and analysis.

Figure 1 illustrated a snippet of large social graph representing relations between four users and three User Generated Contents(UGC). Query which *Find pair of users in a path of friend relationship which user2 has a job and like the documents created by user1* is expressed in SPARQL as:

* This research is supported by the National High Technology Research and Development Program of China (Grant No. 2012AA011002), Natural Science Foundation of China (Grant No. 61300003), Specialized Research Fund for the Doctoral Program of Higher Education(Grant No. 20130001120001).

** Corresponding author.

[1] http://www.foaf-project.org/

[2] http://rdfs.org/sioc/spec/

[3] http://www.w3.org/TR/rdf-sparql-query/

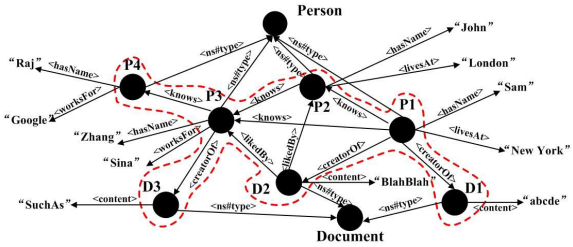

Fig. 1. Examples for a fraction of social network.

```
SELECT DISTINCT ?user1, ?user2 WHERE {
    ?user1 knows* ?user2 .
    ?user1 creatorOf ?doc1 .
    ?user2 worksFor ?organization .
    ?doc1 likedBy ?user2 }
```

For RDF graph in Figure 1, this query returns the result set $R(q) = \{< P1, P3 >\}$. The pattern *?user1 knows* ?user2* states the path consist of zero or more *knows* predicates, it is a property path query pattern.

Challenges. Path queries are of common interest in OSN analysis for discover complex relations among entities. Despite the scalability and flexibility provided by RDF model, Path queries performs poorly and lack of efficient implementation in existing RDF management. Current researches in graph analysis domain are mainly based on pre-constructed reachability indices. Building such indices are both time and memory consuming, especially when dealing with large complex graph. From RDF management point of view, due to costly join for triple pattern matching, there lacks efficient implementation of property path query. Although property path query is recommended by newest SPARQL 1.1 standard, to the best of our knowledge Jena[4], Virtuoso[5] and Sesame[6] are the only three off-the-shelf RDF stores that support standardized path query. They all suffer from performance problems when dealing with path query on large-scale data.

Overview of Our approach. Our approach is motivated by three observations. First, current off-the-shelf RDF store performs well only on join-based star queries. Second, the search space of property path query is restricted to triples representing relations among entities. Third, due to the rich semantics in OSN, most triples are for entities attributes not for relations. Based on these observations, we argue that manage graph topology-related triples into main memory is feasible, and this will greatly enhance the online query performance of property path related sophisticate SPARQL queries. We propose a hybrid framework that has in-memory management of graph topology-related RDF data along with

[4] https://jena.apache.org/
[5] http://virtuoso.openlinksw.com/
[6] http://www.openrdf.org/

disk-based Jena TDB[7] native triple store. In our approach, while loading and indexing triples into TDB, graph topology-related triples are identified and duplicated in main memory. For an online query, we implement a graph traversal based algebra operators for property path pattern, which is more efficient than traditional join-based operator.

Contributions. We summarized our contribution in this paper as:

1. We propose a main memory/disk based hybrid framework for efficient property path query. While leverage the functionality of existing well-established RDF store for BGP query, it specialized in property path pattern query through manage graph topology related data in main memory.
2. We present an algebra operator for property path pattern query. Its realization based on in-memory graph traversal instead of costly join. Using the characteristics of OSN, heuristics for estimating the execution cost of this operator is given that can be used for cost-base optimization.

Organization of the paper. The rest of the paper is organized as follows. Section 2 presents the basic design of our hybrid RDF management framework. Section 3 shows the evaluation results. We conclude in Section 5.

2 Hybrid RDF Data Management

Given vocabulary $\Sigma = V_E \cup V_A \cup E_{EE} \cup E_{EA} \cup L_E \cup L_A$ defined in Table 2, an OSN is represented as a triple set $T_{OSN} = T_G \cup T_A$, where the graph-topology set $T_G \subseteq V_E \times L_E \times V_E$ holds triples representing social entities and relations among them, and the attributes set $T_A \subseteq V_E \times L_A \times V_A$ holds triples representing attributes and theirs relations to social entities.

Table 1. Notations for social graph representation.

Notations	Refers to the set of	Instance of Figure 1
V_E	Nodes for social entities.	$\{P1, D1, \ldots\}$
V_A	Nodes for attributes values.	$\{"John", "London", "abcde", \ldots\}$
E_{EE}	Edges among V_E.	$\{(P1, P2), (P1, D1), \ldots\}$
E_{EA}	Edges between V_E and V_A.	$\{(P1, "Sam"), (D1, "abcde"), \ldots\}$
L_E	Labels for E_{EE}.	$\{< knows >, < likeBy >, \ldots\}$
L_A	Labels for E_{EA}.	$\{< hasName >, < ns\#type >, \ldots\}$

2.1 Architectural Design

In this paper, we focus on efficient implementation of property path query. Consider that path query only related to Triples in T_G (In Figure 1, T_G is represented

[7] `https://svn.apache.org/repos/asf/jena/`

as dashline-encircled part), one direct motivation is that manage T_G in memory will greatly enhance the overall query performance. Based on this motivation, we implement a hybrid RDF data management architecture that manage different query-prone data respectively. This architecture is shown in Figure 2.

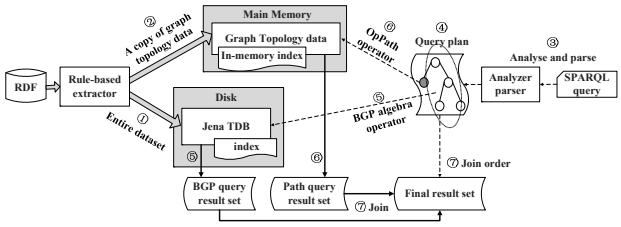

Fig. 2. Hybird RDF management architecture of our approach

Our approach can be thought as a plugin component which override the functionality of corresponding part in TDB. At data loading stage, $\forall t_i \in T_{OSN}$ is loaded into Jena TDB (step ① in Figure 2). At the same time, $\exists t_i \in T_G$ is filtered out based on two kinds of rules:

1. The type of *Objects*. If *Objects* for t_i is a literal, then $t_i \in T_A$.
2. The semantic meaning of *Predicate*. Such as *foaf:knows* defined in FOAF project which state the relation that *Subject User* know *Object User*.

T_G is stored in main memory with subject index(PSO) for forward traversal and object index(POS)for backward traversal(step ②).These indices can be constructed incrementally when topology-related data is extracted and loaded into main memory.

When an online query is submitted, SPARQL parser translates query strings into patterns based on standard abstract syntax tree recommended by W3C (step ③). Analyzer translates these patterns into predefined SPARQL algebra operators and construct execution plan (step ④).An algebra operator generates designated result set from given input (step ⑤). In our approach we implemented a special operator named *OpPath* which only uses in-memory data as input. If a query string in *WHERE* clause is analyzed as property path pattern, *OpPath* operator is added to the query plan. (step ⑥) We explain the design of *OpPath* operator in detail in Section 2.2. Result set of algebra operator is joined to get the final result set. The join order of operators is optimized using cost and selectivity estimation (step ⑦).

2.2 Property Path Algebra Operator

Definition 1 (*OpPath* Operator). *OpPath is a ternary algebra operator that can be defined as OpPath(O,S,P_P). $S, O \subseteq V_{EE}$ can be either bounded or unbounded variables, $|S| = s$, $|O| = o$. P_P is a regular pattern expression defining*

the property path. OpPath(O,S,P_P) operator find existing path from set S to set O, and return all triple sets that each paths is consist of as result set.

Based on research [9] which have testified that graph explorations is extremely efficient and more easy to implement than costly joins, the *OpPath* operator is realized as in-memory Breadth-First Search(BFS). The *OpPath* operator has time complexity $O(|V_E| + |E_{EE}|)$ and space complexity $O(|T_G|)$. It is much less than tradition nested-loop join that has time complexity $O(|V_E| \cdot |E_{EE}|)$.

The cost of *OpPath* operator is the cardinality of result set $R(q)$ for path query pattern q. Existing researches such as G-SPARQL [7] using predefined heuristics which always take $|R(q)|$ as the largest, this is far from optimal. Sparqling Kleene [2] using pre-computed reachability path indices which affects data load efficiency. We consider three factors that affects $|R(q)|$, the average nodes out-degree, the path length l and the pathes that fits for the given pattern. In our approach, we leverages the graph generation model [4] which expects the average out-degree as $d_{out} = |V_{EE}|^{1-\ln c}, 1 < c \leq 2$. We also assumes that nodes in T_G has the same probability of being added to the path, thus the modifying factor of out-degree follows the binomial distribution. For all considerations above, $|R(q)|$ can be approximately estimated as:

$$|R_q| = s \cdot o \cdot \sum_{i=1}^{l} (|V_{EE}|^{(1-\ln c)\cdot i} \cdot (\sum_{i=1}^{l} \frac{l!}{i!(l-i)!} \cdot p^i \cdot (1-p)^{l-i})) \qquad (1)$$

where $p = \frac{|E_{EE}|-|V_{EE}|}{|V_{EE}|}$. $|V_{EE}|$ and $|E_{EE}|$ can be got from the metadata in RDF store. We performs preliminary testing to measure the accuracy of Equation 1 with real all-pair cardinality of dataset in Table 2. For SNIB T_G with average $d_{out} = 12$, $c = 1.75$, $relative\ error = \frac{max(real\ cardinality, estimate\ cardinality)}{min(real\ cardinality, estimate\ cardinality)} - 1$ is about 27%. For DBLP T_G with $d_{out} = 7$, $c = 1.81$, and $error = 32\%$. This preliminary testing shows that the heuristic defined in Equation 1 is with acceptable cardinality estimation error.

3 Evaluation

We used one machine with Debian 7.4 in 64-bit Linux kernel, two Intel Xeon E5-2640 2.0GHz processor and 64 GB RAM for our evaluation. Our approach is compared with Jena(version 2.11.1), Sesame(version 2.7.10) and competitive research G-SPARQL [8]. Jena implements path query based on join while Sesame is based on graph traversal. We also implements our approach with no cost estimation and treats path query as the most costly(denoted as *NoCE*). All evaluations were done 10 times and the results are the averages.

We adopted two datasets, the Social Network Intelligence Benchmark (SNIB)[8] as synthetic dataset, and DBLP as real dataset, all in RDF N-Triples format[9].

[8] Social Network Intelligence Benchmark(SNIB), `http://www.w3.org/wiki/Social_Network_Intelligence_Bench`
[9] `http://www.w3.org/TR/n-triples/`

SNIB dataset is generated using S3G2 [6]. Considering G-SPARQL uses ACM digital library dataset which is not publicly available, we uses the DBLP dataset instead [10]. Statistics of datasets are shown in Table 2. The DBLP dataset has approximately the same characteristics as the *Large Graph Size* experiment in G-SPARQL.

Table 2. Statistics for SNIB and DBLP datasets.

Dataset	Vertices($\|V_{EE}\|$)	Edges ($\|E_{EE}\|$)	Attributes ($\|T_A\|$)	$\|T_G\|/\|T_{OSN}\|$
SNIB	566,472	2,001,333	7,273,177	26%
DBLP	900,440	2,243,827	9,363,166	25%

For offline data loading time and memory expenses, we compare our approach with four competitors, Sesame and Jena in-memory store which store and index data only in memory, Sesame native store and Jena TDB which use disk as triple storage. Results of data loading time is represented in Figure 3(a),and disk usage in Figure 3(b), memory usage in Figure3(c). For our approach only load graph topological data into main memory, it need fewer memory than that of Jena and Sesame in-memory store, but with a little overhead of the data loading time.

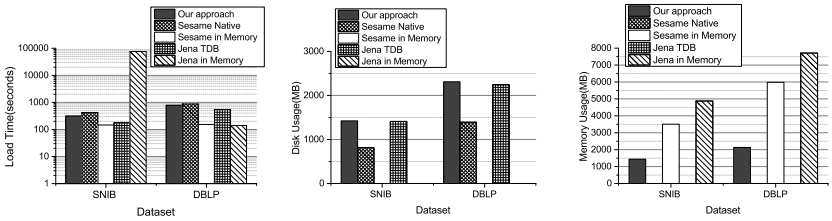

(a) Data loading time. (b) Disk usage. (c) Memory usage.

Fig. 3. *Offline performance evaluation.*

For online query performance, in SNIB benchmark only $Q3$ and $Q5$ are path query related, while for DBLP we chosen 7 out of total 12 queries in G-SPARQL experiments(denoted as Q_g). In order for comparison, we had to rewrite these queries on DBLP dataset. Figure 4(a) shows that our approach achieved the best performance for SNIB $Q3$. As for SNIB $Q5$, the 3-HOP which expressed in UNION clause is explicitly parsed into six joins. This causes an expensive join expenses. Results in Figure 4(b) show that our approach works better than G-SPARQL(got directly from the *Large Graph Size* result in [7]), while *NoCE*

[10] DBLP dataset can be download from `http://sw.deri.org/~aharth/2004/07/dblp/dblp-2006-02-06.rdf`. Same as stated in G-SPARQL, we manually created the co-author relationships between author nodes, which originally recorded as $<$ *creator* $>$ tag in raw dataset.

has approximately the same performance as that of G-SPARQL. This shows that cost estimation for optimal join order can enhance the overall query performance.

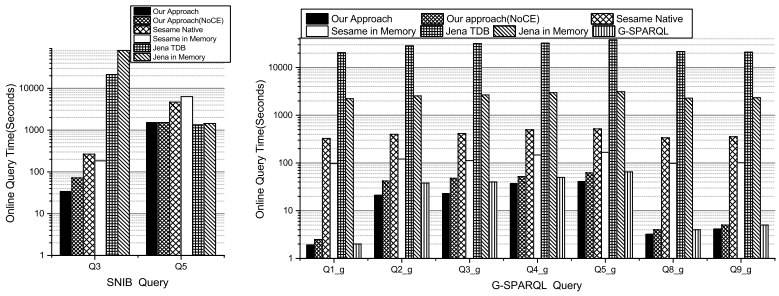

(a) *SNIB query.* (b) *DBLP query performance.*

Fig. 4. *Online Query Performance.*

4 Related Works

Most existing RDF stores uses a relational model to manage data, either in a traditional RDBMS or using a native triple store. They all processes SPARQL queries as sets of join operations using disk-based indices, which are costly for sophisticated joins. Some researches have focused on compressing and managing RDFs in main memory, Trinity RDF [9] is the most prominent among them. It uses graph exploration instead of join operations and greatly boosts SPARQL query performance. But manages all data in memory is not trivial, distributed shared memory increases the complexity for maintenance.

From graph analysis perspective, Property path can be viewed as the label-constraint reachability problem on labeled graph. Though reachability on graph is a further investigated problem, few literatures [1, 10] considered its usage in property graph, which is more common in nature. These researches are mainly focus on building reachability indices in advance, which is time or space consuming and not adequate for large-scale data management.

Besides Jena and Sesame, several frameworks and prototypes have been proposed for path queries [2, 3, 7]. BRAHMS [3] only supports query on paths with predefined length. Sparqling Kleene [2] realizes join based on pre-constructed reachability indices which are space consuming. Our work is mainly motivated by G-SPARQL [7] which use the same hybrid storage and manage graph topological data in memory. Our work different from G-SPARQL in that, first, we are not design a new query language but uses standard SPARQL 1.1 instead, this makes our work more general. Second, G-SPARQL uses index-free pointer-based data structure to representing the graph topological data in memory, in

order to suit for most graph algorithms. Our work only cares about path patterns which BFS algorithm is used to answer such reachability queries. We build simple indices only for facilitating BFS.

5 Conclusion

In this paper we addressed the problem of property path query in RDF data, presented a step towards incorporation of in-memory storage. In our approach we are not trying to invent new wheels, but managed to combine existing effective approaches as well as some technical enhancements. Contrast to traditional RDF management and graph query method, we used in-memory graph traversal instead of costly join to realize path query operator, used simple graph indices other than RDF permutation indices and complex graph reachability indices for efficient graph traversal. Evaluations have shown that our approach is feasible and efficient for process SPARQL property path queries.

References

1. Fan, W., Li, J., Ma, S., Tang, N., Wu, Y.: Adding Regular Expressions to Graph Reachability and Pattern Queries. In: Data Engineering (ICDE), 2011 IEEE 27th International Conference on. pp. 39–50. IEEE (2011)
2. Gubichev, A., Bedathur, S.J., Seufert, S.: Sparqling Kleene: Fast Property Paths in RDF-3X. In: First International Workshop on Graph Data Management Experiences and Systems. p. 14. ACM (2013)
3. Janik, M., Kochut, K.: Brahms: A Workbench RDF Store and High Performance Memory System for Semantic Association Discovery. In: The Semantic Web–ISWC 2005, pp. 431–445. Springer (2005)
4. Leskovec J., Kleinberg J., Faloutsos C.: Graph Evolution: Densification and Shrinking Diameters. ACM Transactions on Knowledge Discovery from Data (TKDD), 1(1) (2007)
5. Neumann, T., Weikum, G.: RDF-3X: A RISC-style Engine for RDF. Proceedings of the VLDB Endowment 1(1), 647–659 (2008)
6. Pham, M.D., Boncz, P., Erling, O.: S3G2: A Scalable Structure-Correlated Social Graph Generator. In: Selected Topics in Performance Evaluation and Benchmarking, pp. 156–172. Springer (2013)
7. Sakr, S., Elnikety, S., He, Y.: G-SPARQL: A Hybrid Engine for Querying Large Attributed Graphs. In: Proceedings of the 21st ACM international conference on Information and knowledge management. pp. 335–344. ACM (2012)
8. Stocker, M., Seaborne, A., Bernstein, A., Kiefer, C., Reynolds, D.: SPARQL Basic Graph Pattern Optimization Using Selectivity Estimation. In: Proceedings of the 17th international conference on World Wide Web. pp. 595–604. ACM (2008)
9. Zeng, K., Yang, J., Wang, H., Shao, B., Wang, Z.: A Distributed Graph Engine for Web Scale RDF Data. In: Proceedings of the 39th international conference on Very Large Data Bases. pp. 265–276. VLDB Endowment (2013)
10. Zou, L., Xu, K., Yu, J.X., Chen, L., Xiao, Y., Zhao, D.: Efficient Processing of Label-constraint Reachability Queries in Large Graphs. Information Systems 40, 47–66 (2014)